NFS / IPv4 + v6
draft-ietf-nfsv4-ipv4v6-00

Go further, faster™

Dhawal Bhagwat
dhawal@netapp.com

Rev: 2010-11-04

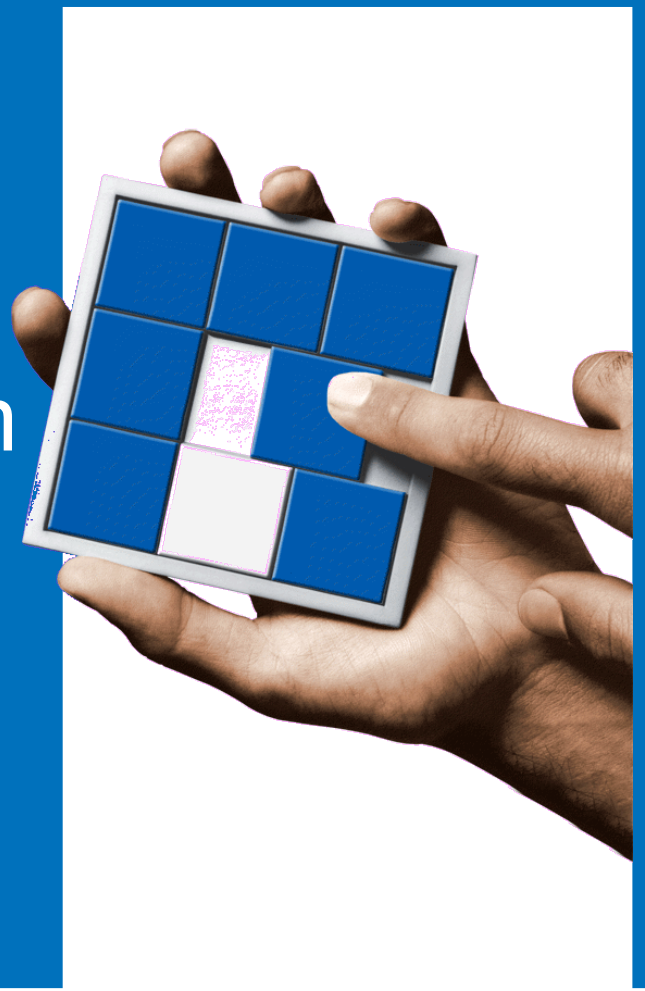

Introduction

- Followup to what was presented at IETF75 by Alex (NetApp).
- This draft addresses issues related to operation of NFS in an Ipv4 + Ipv6 enabled network
  - primarily due to sharing of NFS state (NLM/NSM, and NFSv4) across different protocol address families.

Key points

- RPCBIND / PORTMAP support – MUST use -
  - PORTMAP over IPv4.
  - RPCBINDv3/4 over IPv6.
- NLM/NSM support
  - SHOULD use the "caller_name" (in the NLM_LOCK call), and the "mon_name" (in the SM_NOTIFY call) as the identity of the caller.
  - Using “caller_name” / “mon_name”, perform each action for both IPv4 and IPv6.

Key points (contd.)

- NFSv4 Client Identification
  - client SHOULD use the same client string irrespective of the server address.
    - Relevant for single stack mode too.
- Dual to single stack mode transition
  - Temporary transition – affected states SHOULD be left intact.
  - Permanent transition – affected states for SHOULD be cleared via admin action.

Further course

- NFSv4.x or standalone? - NFSv4 related parts could be included in RFC3530bis; rest standalone.
- Two separate drafts – based on feedback given in IETF75.
- Next steps –
  - Authors will post a new revision for review.
  - Targetting IETF82 for last call.
    - Need members to review.
    - Need WG chair to help reach last call.

Follow up

- ID is available here - http://datatracker.ietf.org/doc/draft-ietf-nfsv4-ipv4v6/
- Comments - nfs4@ietf.org or dhawal@netapp.com